CLAIMS

1. A method for configuring a computing system having a plurality of tunable parameters, the method comprising the steps of:

formulating an optimal setting for the configuration of at least two of the plurality of tunable parameters as a black box optimization problem; and

solving the black box optimization problem using a smart hill climbing method, where the optimal setting yields improved computing system performance.

2. The method of claim 1, wherein the smart hill climbing comprises:

performing a global search; and

performing a local search on a starting point identified by said global search.

3. The method of claim 2, wherein the starting point represents one possible solution to the black box optimization problem.

4. The method of claim 2, wherein the local searching step searches in a neighborhood of the starting point for a solution that will yield better computing system performance than the starting point.

5. The method of claim 4, wherein the local searching step applies a gradient-based sampling method to search the neighborhood.

6. The method of claim 2, wherein at least one of the global and local searching steps comprises a Latin Hypercube Sampling method.

7. The method of claim 6, wherein the Latin Hypercube Sampling method is weighted.

8. The method of claim 7, wherein the weighted Latin Hypercube Sampling method is adapted to generate a plurality of random configurations of a specified dimension.

9. The method of claim 8, wherein the weighted Latin Hypercube Sampling method comprises the steps of:

generating a plurality of permutations of the random configurations, wherein the number of permutations corresponds to the number of dimensions;

dividing a parameter range of each dimension into a plurality of non-overlapping intervals with equal probabilities, wherein the number of intervals corresponds to the number of random configurations; and

generating a random sample within each interval of each dimension.

10. The method of claim 2, wherein at least one of the global and local searching steps comprises the steps of:

establishing a local search neighborhood around an initial candidate point;

generating a plurality of samples from within the local search neighborhood;

updating information regarding optimal computing system parameter configuration, based on the plurality of samples generated; and

processing the plurality of samples to generate a second candidate point for the solution to the black box optimization problem.

11. The method of claim 10, wherein the step of establishing the local search neighborhood comprises:

generating an initial sample;

identifying an initial candidate point that will yield the best computing system performance; and

setting the initial candidate point as the center of a neighborhood of a predefined size.

12. The method of claim 11, wherein the initial sample is generated using a weighted Latin Hypercube Sampling method.

13. The method of claim 11, further comprising the steps of:

modifying the local search neighborhood so that the second candidate point is at the center of the local search neighborhood, if the next candidate point will yield better computing system performance than the initial candidate point; and

generating a subsequent sample; and

updating information regarding optimal computing system parameter configuration, based on the subsequent sample.

14. The method of claim 10, wherein the step of processing the plurality of samples further comprises the steps of:

constructing a quadratic fitting curve based on the generated samples, wherein a quadratic fitting curve is constructed for each dimension;

identifying a minimal point according to the quadratic fitting curve; and

combining the minimal points for all dimensions to generate the second candidate point.

15. The method of claim 14, further comprising the steps of:

constructing a new quadratic fitting curve including the second candidate point, if the second candidate point will not yield better computing system performance than the initial candidate point; and

generating a subsequent candidate point based on the new quadratic fitting curve.

16. The method of claim 15, further comprising the steps of:

modifying the local search neighborhood so that the subsequent candidate point is at the center of the local search neighborhood, if the subsequent candidate point will yield better computing system performance than the initial candidate point; and

generating a subsequent sample; and

updating information regarding optimal computing system parameter configuration, based on the subsequent sample.

17. The method of claim 15, further comprising the step of:

shrinking the size of the local searching neighborhood, if the subsequent candidate point will not yield better computing system performance than the initial candidate point.

18. The method of claim 17, further comprising the step of:

generating a new sample from the local searching neighborhood, if the size of the local searching neighborhood has not shrunk beyond a predefined threshold; and

updating information regarding optimal computing system parameter configuration, based on the new sample.

19. The method of claim 1, further comprising the step of:

applying said optimal setting to configure an application server.

20. A computer readable medium containing an executable program for configuring a computing system having a plurality of tunable parameters, where the program performs the steps of:

formulating an optimal setting for the configuration of at least two of the plurality of tunable parameters as a black box optimization problem; and

solving the black box optimization problem using a smart hill climbing method, where the optimal setting yields improved computing system performance.

21. The computer readable medium of claim 20, wherein the smart hill climbing comprises:

performing a global search; and

performing a local search on a starting point identified by said global search.

22. The computer readable medium of claim 21, wherein the starting point represents one possible solution to the black box optimization problem.

23. The computer readable medium of claim 21, wherein the local searching step searches in a neighborhood of the starting point for a solution that will yield better computing system performance than the starting point.

24. The computer readable medium of claim 23, wherein the local searching step applies a gradient-based sampling method to search the neighborhood.

25. The computer readable medium of claim 21, wherein at least one of the global and local searching steps comprises a Latin Hypercube Sampling method.

26. The computer readable medium of claim 25, wherein the Latin Hypercube Sampling method is weighted.

27. The computer readable medium of claim 26, wherein the weighted Latin Hypercube Sampling method is adapted to generate a plurality of random configurations of a specified dimension.

28. The computer readable medium of claim 27, wherein the weighted Latin Hypercube Sampling method comprises the steps of:

generating a plurality of permutations of the random configurations, wherein the number of permutations corresponds to the number of dimensions;

dividing a parameter range of each dimension into a plurality of non-overlapping intervals with equal probabilities, wherein the number of intervals corresponds to the number of random configurations; and

generating a random sample within each interval of each dimension.

29. The computer readable medium of claim 21, wherein at least one of the global and local searching steps comprises the steps of:

establishing a local search neighborhood around an initial candidate point;

generating a plurality of samples from within the local search neighborhood;

updating information regarding optimal computing system parameter configuration, based on the plurality of samples generated; and

processing the plurality of samples to generate a second candidate point for the solution to the black box optimization problem.

30. The computer readable medium of claim 29, wherein the step of establishing the local search neighborhood comprises:

generating an initial sample;

identifying an initial candidate point that will yield the best computing system performance; and

setting the initial candidate point as the center of a neighborhood of a predefined size.

31. The computer readable medium of claim 30, wherein the initial sample is generated using a weighted Latin Hypercube Sampling method.

32. The computer readable medium of claim 30, further comprising the steps of:

modifying the local search neighborhood so that the second candidate point is at the center of the local search neighborhood, if the next candidate point will yield better computing system performance than the initial candidate point; and

generating a subsequent sample; and

updating information regarding optimal computing system parameter configuration, based on the subsequent sample.

33. The computer readable medium of claim 29, wherein the step of processing the plurality of samples further comprises the steps of:

constructing a quadratic fitting curve based on the generated samples, wherein a quadratic fitting curve is constructed for each dimension;

identifying a minimal point according to the quadratic fitting curve; and

combining the minimal points for all dimensions to generate the second candidate point.

34. The computer readable medium of claim 33, further comprising the steps of:

constructing a new quadratic fitting curve including the second candidate point, if the second candidate point will not yield better computing system performance than the initial candidate point; and

generating a subsequent candidate point based on the new quadratic fitting curve.

35. The computer readable medium of claim 34, further comprising the steps of:

modifying the local search neighborhood so that the subsequent candidate point is at the center of the local search neighborhood, if the subsequent candidate point will yield better computing system performance than the initial candidate point; and

generating a subsequent sample; and

updating information regarding optimal computing system parameter configuration, based on the subsequent sample.

36. The computer readable medium of claim 34, further comprising the step of:

shrinking the size of the local searching neighborhood, if the subsequent candidate point will not yield better computing system performance than the initial candidate point.

37. The computer readable medium of claim 36, further comprising the step of:

generating a new sample from the local searching neighborhood, if the size of the local searching neighborhood has not shrunk beyond a predefined threshold; and updating information regarding optimal computing system parameter configuration, based on the new sample.

38. The computer readable medium of claim 20, wherein the executable program further comprises the step of:

applying said optimal setting to configure an application server.

39. Apparatus for configuring a computing system, comprising:

means for formulating an optimal setting for the configuration of at least two of the plurality of tunable parameters as a black box optimization problem; and

means for solving the black box optimization problem using a smart hill climbing method, where the optimal setting yields improved computing system performance.

40. The apparatus of claim 39, further comprising:

means for applying said optimal setting to configure an application server.